

One problem in any technical field is that technical fields have technical terms which sound self-explanatory and are not. I am sure the world is full of people who believe that they understand the theory of relativity, except for the mathematical details. After all the theory of relativity says things are relative, I understand that. But that of course is not what the theory of relativity says. Similarly in economics the term market failure sounds self-explanatory, but I will be arguing today that it is a deceptive label, that correctly understood not all ways in which markets can fail are market failure and market failure exists in many context other than markets. That it is a problem that humans face in markets, in political institutions, in armies, in families, in firms, and I will try to offer some examples.

I will try to show both why market failure really is an argument against laissez faire, an argument against free markets, but why it is, I think, an even stronger argument against the alternative to free markets. Let me start with a very simple example of market failure. I want you to imagine that it is a little less than a thousand years ago and you are one of a line of about 5000 men with pikes, and you're pointing the pikes in that direction. The reason you're pointing the pikes in that direction is that a similar number of men on horseback with lances are coming at you. Since this is Switzerland I suppose they are flying the Habsburg colors. You do a very quick cost benefit calculation. You say if all of us stand, and if we plant our pikes and hold them well, with luck we can break their charge, and if we break their charge some of us will die but most of will live. If we try to run, horses run faster than we do. So I should stand.

Nope, sorry, I made a mistake; I said "we." I don't control him and him; the question is what should I do. If everybody else stands and I run, since there are 5000 men in this line one man running will have very little effect on the chance that we stop the charge, and I won't be one of the ones killed. If everybody else runs my only chance is to run first and fastest. So whatever the rest of the people do I'm better off running. Everybody makes that calculation, we all run and most of die.

Welcome to the dark side of rationality.

What this is illustrating is a situation where if each individual makes the individually rational correct choice everyone is worse off. That, I think, is the correct understanding of market failure. Market failure describes those situations where individual rationality does not lead to group rationality.

There are lots of other familiar examples. Those of you who have studied either economics or game theory are probably familiar with the prisoner's dilemma, which is a two person version of this situation. Two men jointly commit a crime, they're arrested, the prosecutor tells each of them if you confess I'll go somewhat easy on you, if he doesn't confess and you do I'll go really easy on you because I'll need your confession, if neither of you confesses I can still convict you of something, if both of you confess I can convict you both of something much more serious, and if he confesses and you don't I'll give you the maximum penalty. The prosecutor can arrange the threats he makes in such a way that each of the two prisoners is better off confessing than not, because if the other

person doesn't confess then his confession will get him a very light sentence and if the other person does confess and he doesn't he'll be treated very badly. Yet, they would both be better off staying silent.

There are lots of other examples; the London fogs of the 19th Century were due to people burning coal in their fireplaces. It's possible that burning something else would have been better, but each individual correctly saw that coal was the best fuel for him and he bore only a small part of the cost of the pollution from the fog.

One point worth noting, however, is that whether or not you're in a situation of market failure partly depends on the particular rules of the games you are playing. That suggest that sometimes, if you're designing laws or for that matter setting up a firm, you can engineer around market failure. You can set things up to make individual rationality lead to group rationality.

My favorite example of this is the camel puzzle. Two Bedouins are riding through the desert complaining about how slow their camels are. The first Bedouin says "this beast of mine is the slowest beast in all the deserts of Arabia; it's going to be an hour before I reach that oasis." The other one says "you think your camel is slow? Let me tell you about my camel."

So they get into an argument about which one has the slower camel, and eventually they agree on a bet. Whichever one's camel gets to the oasis last his owner wins the bet and the other will pay him a golden dinar. You can see what happens; one starts moving slowly and the other moves more slowly and the other more slowly still, because both of them want to win the bet. An hour later there are two Bedouins sitting their camels stock still in the blazing sun a mile away from the oasis.

At this point a wise man comes by and the wise man says "what are you two idiots doing sitting your camels in the blazing sun a mile from the oasis?" and they get off their camels to explain the problem. The wise man says two words to them and they leap back on the camels and race off for the oasis as fast as ever they can.

What are the two words? "Switch camels." Each one gets on the other's camel and now he wants to go fast and not slow. That's a nice example of where a very small change results in engineering around the market failure.

How does market failure come into economic and political arguments? Economists tend to assume individual rationality. In fact, one definition of economics, the one I prefer, is that it is the approach to understanding behavior that starts with the assumption that individuals have objectives and tend to choose the correct way of achieving them. So you might think that if you assume rationality then the case for laissez faire is perfectly clear; everybody rationally does what's in his own interest and are all better off as a result. Market failure demonstrates that that argument is not true. Under some circumstances everybody acting in his own interests gives you a desirable result, that's Adam Smith's famous invisible hand. But there are other circumstances where

everybody acting in his own interests gives an undesirable result—and that looks like an argument for government making people do something else.

There are a number standard economic textbook categories of market failure; one of them is the public good problem. Suppose I am producing something and can't control who gets it, a radio broadcast for instance—if it's produced it will be available to all of the members of a pre-existing group, all the people within range of my transmitter. How do I pay for it? For producing ordinary things I can say to each consumer "you don't get it unless you pay," but for a radio broadcast I can't do that. You could think of a lot of other examples of a public good, such as a flood control dam or national defense for Switzerland, lots of examples where you have something desirable which if it's produced a bunch of people will get it and you can't control who. That's a public good, and the public good problem is the failure to produce things that are worth the cost of producing. Where the value to those benefitted is larger than the cost, but they don't get produced because there's no way the producer can pay for it. That is one example of a form of market failure.

A very closely related one is what economists called externalities, cases such as air pollution where my action results in some costs and some benefits for me, but also costs for other people. I take into account the costs and benefits for me, but not for other people, so if my benefits are larger than my costs I do it even though it may be that our benefits are less than our costs. That's externalities.

It's worth noticing that some market failure problems can be and are engineered around. Sometimes on the ordinary private market ingenious people can find ways of solving these problems. I like to say that another name for market failure is a profit opportunity, because market failure means there is something not happening which, if it did happen, would produce net benefits. If you can figure out a way of making it happen and collecting some of those benefits you might be able to make a profit—but not always.

Radio programs are a good example of a solution. A radio broadcast is a pure public good, it's obviously impossible for it to be produced privately and yet they are produced privately. That's because some unknown genius thought up the idea of producing not one public good but two. One public good with a positive cost of product and a positive value to the consumer, which is called a radio program, and one public good with a negative value to the consumer and a negative cost of production called an advertisement. An advertisement has a negative cost of production because people will pay the radio station to put it on. You tie the two goods tightly together and give away the bundle; that's how that particular problem was solved. There are other examples. Some externality problems can be solved by putting multiple activities under one roof. If you think of a housing development or a shopping center the owner of the shopping center can provide convenient parking, which is a positive externality for the stores, because it gets them more customers, and he's collecting in the rent for the stores enough so he's willing to do that.

If I had more time I could go through other examples where people successfully solve such problems. But there is no good reason to suppose that they will all be solved; there is no good reason to suppose that in a laissez faire society all goods worth producing will get produced or that no goods not worth producing will get produced. In that sense market failure really is a legitimate criticism of a laissez faire society, a legitimate argument for government intervention. To put it differently, market failure implies that laissez faire will get us a less good result than we would get if we had regulation by a perfectly wise and perfectly benevolent regulator.

We have, however, a striking shortage of perfectly wise and perfectly benevolent regulators. And it turns out that if we consistently apply the same standards to the alternatives of laissez faire, market failure is also a reason why political institutions don't give us the results we would like them to give us.

Let me start with a simple example, what I like to think of as the civics class model of democracy. The simple school boy model of why democracy is supposed to work. The idea is that if politicians do good things you vote for them, if they do bad things you vote against them, so politicians have to do good things.

The problem with that is that in order for it to work the voters have to know whether the politician has done good things. I don't live in Switzerland, and perhaps you have more forthcoming and transparent politicians than we do, but as far as I know no politician in America has ever run on the campaign slogan "I'm the bad guy." I don't think anyone has ever introduced a bill into Congress which was labelled "a bill to make farmers richer and city folks poorer." Every year bills with that purpose are introduced into Congress, bills that cost many billions of dollars, because the main purpose of U.S. farm policy is to raise the price of agricultural products, which makes farmers richer and everybody else who consumes those products poorer.

But that's not how they label them. They always explain how it is necessary for the national welfare that we maintain a healthy farm sector in order that we not all starve to death, and therefore they're doing it for everybody. The same pattern is true for all political policies. That means that in order for the civics class model to work the individual voter has to devote a lot of time and energy to figuring out both what his representative is doing and what he ought to be doing. What is the incentive to devote time and energy to that? The answer is there is none, because the voter in a large democracy, in the U.S., surely in Switzerland at the national level, probably at the canton level, perhaps not at the local level, knows that the chance that his vote will change the outcome is very, very close to zero. My estimate is about one in 5,000,000 or 10,000,000 for a U.S. presidential election.

It is not worth bearing large costs for one chance in 5,000,000 of improving the outcome for you. It's true that if I elect the right guy, and if the right guy is the right guy not just for me but for the country, many other people will benefit. But that's a public good; I don't get to collect from their benefit. So being an informed voter involves producing a public good for a very large public, public goods are under produced, typically are not

produced, and the result is what is referred to in the economic analysis of politics as rational ignorance. It is rational to be ignorant when the information cost more than it is worth, and it is almost always the case that political information, at least for the individual voter, costs more than it is worth. So that describes one case in which market failure undercuts the fundamental mechanism that's supposed to make our government's work.

Let me give you a different example of a different sort of market failure, and this has to do with the question of to what degree future costs and benefits are taken account of in decisions. Let me start with the market. It would seem at first that nobody would invest in a project whose payoff took more than his lifetime, but it's not true.

Suppose there is some valuable hardwood, let's say black walnut which is a very pretty wood, which is slow growing, and I'm 60 years old and it takes 50 years to bring a field full of black walnut trees ready to harvest. You might say why would I plant them, and the answer is the reason I plant them is that 20 years from now I can sell them to somebody else who can then wait 30 years to harvest them or maybe 20 and then sell them to someone else. Black walnut trees 20 years along are that much closer to maturity. Hence in a private market it pays to allow for costs and benefits, so far as you can predict them, into the distant future.

This depends, however, on secure property rights. Suppose I know that each year there's a significant chance that a private thief will drive onto my field at night, cut down my partly grown trees and sell them, or a significant chance that a public thief will announce that black walnut trees are a national resource and confiscate them. I will then only plant the trees if the return is so large it allows for the risk of losing them year by year. So what that suggests is that the more insecure property rights are the more rationally short-sighted you are going to be, the less willing to sacrifice present for future. Politicians in a democratic system have very insecure property rights. As I used to put it in a previous administration, Bill Clinton could rent out the White House but he couldn't sell it.

Suppose you're a politician and you see that there is something you could do which will lose you votes this year, making it politically expensive, but will produce a large benefit forty years later. Forty years from now you're not going to be President. Fifty/fifty chance the President will even be the other party from yours. There is no mechanism equivalent to selling those black walnut trees to give you an incentive to plan that far ahead. So the result is that although politicians talk a lot about concerns with the future, all the real attention is on what the unemployment and inflation rate will be at the next election. That's sensible, rational behavior by politicians; they know they don't have secure property rights and so they are rationally short-sighted. That's another example of market failure.

Let me go on to a particular example which provides something better than the civics class model for understanding democracy. That's the market for legislation.

Suppose that you are part of an interest group, say the auto industry, and you would like an auto tariff because that will raise the price of automobiles, which will let you make more money. An auto tariff is a public good for a public consisting of the auto companies plus the auto workers; if two of the auto companies bribe politicians to pass a tariff, the auto company that didn't contribute will still get the benefit of the tariff. But public goods are much easier to produce for very small publics; if you're talking about four or five companies and one union they can get together, they can say to each other "here's what I'll chip in if you chip in some, and if you refuse it's not going to happen." That kind of unanimous contracting is one of the solutions for the public good problem for small publics.

You would expect in general that the smaller and more concentrated, the more organized, a group is the easier it is for it to produce a public good for that group. Passing an auto tariff is a public good for a fairly small group, stopping the auto tariff is a public good for a very large group, since the people who lose from the tariff include everybody who buys automobiles and everybody who produces export goods—if we import less we're going to export less. If the effect of the auto tariff is a \$10 billion gain for the auto companies and a \$20 billion loss for everybody else, the auto companies will raise \$2 billion of that \$10 billion and offer it to politicians to buy the legislation. The victims will raise \$10,000 of that \$20 billion and offer it to politicians to block the legislation and the tariff will pass. This explains why, although the economics of trade was worked out about 200 years ago by David Ricardo (one of my favorite economists for a number of reasons), nonetheless from then until now there have only been a very small number of countries that have followed the logic of that analysis and actually had free trade.

Generalizing the argument I've made, we can expect that the usual tendency on the political market is going to be to spend resources transferring from dispersed interests to concentrated interests. I think this is what we pretty consistently observe. Gary Becker, a fairly prominent economist, has an interesting observation about farm policy around the world based on this theory that seems to fit the data. The world is divided into two kinds of countries. There are countries where farmers are a small minority like the U.S. and Europe; in those countries the function of farm policy is to enrich farmers at the cost of everybody else. There are countries where the farmers are the large majority like a lot of the Third World, and there the function of farm policy is to make the dispersed peasants poorer in order to have cheap food to bribe the support of the urban mob, which is politically more organized and more influential.

I believe, by the way, there is one exception. As far as I can tell New Zealand is the only country in the world that neither especially taxes nor especially subsidizes farmers. I don't know if this is a miracle or just experimental error, but on the whole Becker's observations seems to be correct. I can go onto other examples. One of the problems, which actually I suppose is relevant to an important Swiss industry, is regulation of medical drugs. The fundamental problem is that if you approve a drug that then turns out to have bad side effects that fact will be very visible and the regulator who approves it will lose his job. If you stall a drug which does not have bad side effects for a few years the actual result will be that people will die who would otherwise have lived, but nobody

will know who; they'll be statistical deaths rather than headlines. The result, given information costs, is that there is a strong incentive to over regulate.

Let me give you one more example, and this is not from the regulatory or legislative side, but from the judicial side. There is a particular law case in the U.S. history, not that long ago, I suppose maybe 30 years ago or so, a little more, in which a panel of judges made a mistake which I believe an intelligent high school student should be ashamed of. That mistake almost certainly killed thousands of people, and none of those judges ever suffered any penalty for that, because judges are not liable for mistakes. This was a polio vaccine case involving the old live polio vaccine before they had the modern medicine. It was unavoidably hazardous; there was about one chance in a million that vaccination would give you polio. The company that produced it, which I think was Merck, informed the doctors and nurses who gave the vaccination about this risk but did not make sure that they informed the people who were being vaccinated.

Somebody got vaccinated, and I think...I'm not sure if he got polio or if his kid was vaccinated and he got polio, but somebody got polio as a result of vaccination, and sued. The court's position was that the vaccine was unavoidably hazardous, so it wasn't the fault of Merck that the risk existed. The question was whether they should have made sure people were warned. The argument that the court came up with was whether they were obliged to make sure that the people who were vaccinated were warned depended on whether the danger was enough so that somebody might reasonably decided not to get vaccinated because of it. So the court compared the risk of polio due to vaccination to the risk of polio if you weren't vaccinated. They compared the one in a million chance of polio from the vaccination with the roughly one in a million annual risk of getting polio. The vaccination lasts for a lifetime. They were making mistake of about a factor of 40 by using the risk of getting polio in one year instead of the risk of getting polio in a lifetime. As a result of making that mistake they concluded that Merck should have made sure that the people who got it were warned, that Merck was liable for a whole of lot money.

The result of this was, for I think a couple of years, to sharply reduce the introduction of new vaccines. Eventually, Congress immunized the drug companies in order to get vaccines back. I don't know how many people died because of that delay; it's hard to believe it was less than a few thousand. As far as I know none of those judges ever owed any liability to anybody, because judges are not responsible for their mistakes in our legal system.

Market failure comes because somebody makes a decision where other people bear the costs or get the benefits. That's true of all judges, it's true of all legislators, it is true essentially of all regulators. Almost everything government does has the characteristic that the person making the decision does not himself bear the cost that results. It's true of voters too. What that suggests is that while market failure does indeed exist on private markets, it is the exception on private markets and the rule on the public market. So one should view market failure as a reason why the laissez faire works imperfectly, and a reason why the political alternative works catastrophically badly.

One of the things that explains is why market failure is a better predictor for political outcomes than for private outcomes. Market failure theory tells us that there are no private radio broadcasts, and yet there are. Market failure theory tells us that farm policy should have the pattern that we've described and it does. Market failure better explains the political market than the private market.

Let me close by pointing out some of the other contexts in ordinary life where market failure matters. For one example, assume we have two people who are sharing living quarters; whether they're married or lovers or roommates, it doesn't really matter for my argument. But both of them cook—that's the critical assumption—and they believe in equality so they share the work.

One possibility is one night I cook and I clean up, the next night you cook and you clean up. The other is that one night I cook and you clean up and the next night you cook and I clean up. The question is which makes more sense? I only want answers from people who have actually cooked—they have the technological information relevant to answer the question. Which is the right answer, anybody want to tell me?

Clean up your own mess, that's right. The right answer is that how much mess you make is not fixed, how much mess you make in cooking is the result of a whole bunch of decisions you make. Are you going to cook a meal that can be done in one pot or seven? While the water is coming to a boil are you washing out the frying pan that you used for the previous stage of the cooking, and so forth and so on? If each person does his own cleaning up then he has an incentive to hold down the amount of mess he makes whenever it's not too costly to do so; if the other person is going to do the cleaning up, he doesn't.

That also explains the apparently very foolish rule that parents impose on small children that they have to clean up their own mess. Most of the time, in the case of small children, the parent can clean up the mess much more easily than the child, and the child cleaning up the mess might well make more mess. But establishing the principle means that the child has an incentive not to make messes.

Let me go on to give a different example from my profession, from teaching. This is what I like to refer to as the silent student problem. I give a perfectly clear explanation of some point in economics, and I then ask "did everybody follow that," and nobody says anything. So I keep going, and only discover my mistake when I grade the final exam and discover that almost nobody understood that point. What's happening? Each student feels that if he says he didn't understand it he will look stupid to the other students, he might look stupid to me which might lower his grade, and any benefit from getting me to repeat the explanation will be shared with all the rest of the students. That's a public good.

I have a solution. My solution is to equip each classroom with a small button at each place which can be unobtrusively pushed with the student's foot. At the back of the classroom is a sign showing how many buttons are pushed. So I come to the appropriate

place in the lecture and ask everybody who clearly followed my explanation to push his button. The number two appears on the screen at the back of the room and I go back and explain it again. That's again a case of engineering around the public good problem.

Finally, consider the application to running firms, something I have not done but some of you may have. One useful way of thinking about what you are doing is that you're trying to design a set of rules that as nearly as possible avoid these problems, rules set up so that as nearly as possible each person gets the benefits and pays the costs of the decisions he himself makes.

Notice that throughout this discussion I am consistently applying the same assumptions to everybody. I don't assume that politicians are either more or less benevolent than businessmen, I don't assume that roommates are more or less benevolent than politicians and businessmen, I don't assume that voters are different, I don't assume that politicians are stupid, that voters are stupid. I'm willing to assume that everybody is rational and everybody is simply acting in his rational self-interest in the same way. Self-interest does not mean selfish in the narrow sense. Self-interest only means that each person has a set of objectives, some of which might be narrowly selfish, some of which might be entirely benevolent. But each person is pursuing his own objectives, people have different objectives and therefore you have to take account of the fact that each person is acting to achieve his own objectives.

Thank you.
[End of Audio]